

White Paper

What's Needed for Cloud Computing?

Focus on Networking and WAN Optimization

By Jon Oltsik

June, 2010

This ESG White Paper was commissioned by Riverbed Technology and is distributed under license from ESG.

Contents

Executive Summary	3
De-mystifying Cloud Computing.....	3
The Cloud Consumption Model.....	3
The Cloud Computing Services Model.....	4
Cloud Computing Deployment Models	5
The Evolution of Cloud Computing.....	5
The Cloud Depends Upon the Network.....	6
WAN Optimization is a Critical Underpinning of Cloud Computing	7
Will Cloud Computing Change WAN Optimization?	7
The Bigger Truth	9

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.

Executive Summary

Just what is cloud computing anyway? Skeptics might say it is nothing but industry hyperbole, visionaries might say it is the future of IT. In reality, both statements are true—cloud computing has been embellished by the tech industry, but it does hold real potential for new types of on-demand dynamic IT services. This paper seeks to clarify the definition of cloud computing, identify how far along users are in terms of cloud deployment, and examine the role of the network in the cloud computing model. In summary:

- **Cloud computing spans numerous models.** Perhaps the most confusing thing about cloud computing is that it isn't one thing; rather, it can be a number of customizable options to fit IT and business requirements. This paper will look at various alternatives around the cloud computing consumption, services, and deployment models.
- **Large organizations are on an evolutionary journey to the cloud.** Cloud computing builds upon current IT trends like data center consolidation and server virtualization. Enterprises will follow a pragmatic path to the cloud by adopting new technologies, transitioning from physical to virtual IT assets, and adapting existing IT best practices to a new dynamic world.
- **The network acts as the foundation for cloud computing.** Cloud computing moves Web-based applications to the Internet inexorably tying user connectivity and productivity to networking equipment. Of all networking technologies available, WAN optimization will play a major role in the transition to the cloud. As this happens, WAN optimization must become a virtual service, support mobile users, support innovative applications and protocols, and provide network visibility at a granular level.

De-mystifying Cloud Computing

Pick up any technology magazine or go to any IT trade show; regardless of where you look, everyone is talking about cloud computing. Of course, all of this hype has resulted in tremendous confusion. If you ask ten IT professionals to define cloud computing, you may end up with ten different responses.

Clearly, there is a need for a baseline definition and a discussion of the many aspects comprising cloud computing. ESG believes that a detailed characterization of cloud computing must include three things:

1. **The consumption model.** Cloud computing offers a unique way to consume compute, network, and storage resources.
2. **The services model.** Cloud computing can be used as a platform for different activities up and down the technology stack.
3. **The deployment model.** Cloud computing is not a one-size-fits-all platform. Rather, there are numerous ways to deploy and utilize clouds.

The Cloud Consumption Model

At the most fundamental level, cloud computing provides flexible real-time access to a shared pool of computing resources (e.g., networks, servers, storage, applications, and services). Indeed, one of the main attractions of cloud computing is its capability to provide on-demand IT resources and services offering rapid provision and de-provisioning as well as “pay by the drink” pricing.

In its definition of cloud computing, The National Institute of Standards and Technology (NIST) describes cloud computing as having the following essential characteristics:¹

- **On-demand self-service.** A consumer can unilaterally provision computing capabilities such as server time and network storage as needed. This can happen automatically, without human interaction, system administration, or service provider support.

¹ Source, NIST, Computer Security Division, Computer Security Resource Center, NIST Definition of Cloud Computing v15.

- **Broad network access.** Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).
- **Resource pooling.** The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources, but may be able to specify location at a higher level of abstraction (e.g., country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.
- **Rapid elasticity.** Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and then rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.
- **Measured service.** Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized service.

From an enterprise IT perspective, the overwhelming benefit with cloud computing is flexible on-demand access to IT resources without the usual purchasing, deployment and management overhead. Need storage resources for the next 2 weeks? With cloud computing, terabytes of storage can be available instantly, with the swipe of a credit card. After the two week period, customers pay for time and capacity.

The Cloud Computing Services Model

Cloud-based IT resources can be consumed for various reasons in myriad ways. That said, the IT industry has really rallied around three standard cloud computing services models:

1. **Infrastructure-as-a-service (IaaS).** This is really the most basic cloud service model, aligning the on-demand resources of the cloud with tactical IT needs. In this way, IaaS is similar to managed services offerings of the Internet era (i.e., hosting services, storage service providers (SSPs), etc.). The primary difference is that cloud resources are virtual rather than physical and can be consumed on an as-needed basis. In other words, enterprise consumers pay for virtual machines (VMs), storage capacity, and network bandwidth for a variable amount of time rather than servers, storage arrays, and switches/routers on a contractual basis. Typically, enterprises have no control over the underlying cloud technology. As described above, IaaS can be utilized as a temporary resource or used for years at a time. IaaS prices are based upon two factors: IaaS resource consumption and the duration of use.
2. **Platform-as-a-service (PaaS).** Simply stated, PaaS provides the capability to build or deploy applications on top of IaaS. Typically, a cloud computing provider offers multiple application components that align with specific development models and programming tools. For the most part, PaaS offerings are built upon either a Microsoft-based stack (i.e., Windows, .NET, IIS, SQL Server, etc.) or an open source-based stack (i.e., the "LAMP" stack containing Linux, Apache, MySQL, and PHP).
3. **Software-as-a-service (SaaS).** In this model, an entire business or set of IT applications runs in the cloud. Enterprise consumers outsource the entire underlying technology infrastructure to a SaaS provider and thus have no responsibility or management oversight for SaaS-based IT components. Users typically access these applications over the Internet through a thin client interface such as a Web browser. SaaS examples include Gmail from Google, Microsoft "live" offerings, and salesforce.com.

There is no hierarchy in these service offerings; rather, CIOs can choose any or all cloud service offerings that fit their needs.

Cloud Computing Deployment Models


Just like the cloud services models, cloud computing can be deployed in a number of ways depending upon factors like security requirements, IT skills, and network access. The IT industry has outlined four cloud computing deployment models:

1. **Private cloud.** The cloud infrastructure is operated within a single organization. In this case, internal groups such as business units consume resources and services provided by a single internal (i.e., the IT department) or external cloud computing provider.
2. **Community cloud.** A community cloud is a superset of a private cloud. The cloud supports the needs of several or an extended community of organizations. Again, community clouds can be built and operated by members of the community or third-party providers.
3. **Public cloud.** The cloud infrastructure and services are available to the general public. Examples of public clouds include [Amazon](#) Elastic Compute Cloud (EC2), [Google](#) App Engine, [Microsoft](#) Azure, or [Terremark](#) Cloud Computing services.
4. **Hybrid cloud.** The cloud infrastructure amalgamates private or community clouds with public clouds. In this case, private or community cloud services have the capability to extend or “burst” to consume public cloud resources.

The Evolution of Cloud Computing

Enterprise organizations will likely experiment with cloud computing, carefully choosing projects that benefit from cloud’s features and cost benefits as they develop more formal cloud computing strategies. This evolution has already begun as enterprise businesses take a “crawl, walk, run” approach that builds toward an eventual cloud implementation. Figure 1 depicts an IT progression that begins with basic use of server virtualization, evolves into responsive and dynamic internal IT, and finally advances to cloud computing.

Figure 1. Cloud Computing Migration Phases



Source: Enterprise Strategy Group, 2010

The phases of the model include:

- **Test and development.** This phase introduces cloud for proof-of-concept use. During this initial phase, IT becomes comfortable with server virtualization and gains experience with system performance, application response times, and technology stability.
- **Consolidation.** This phase is highlighted by the migration of physical servers to virtual machines—typically referred to as P2V. At this point, IT rapidly moves workloads that have been identified as viable candidates and gives them the green light for production usage on the virtualized infrastructure.
- **Enterprise.** This phase is a significant milestone where the business chooses a virtualization platform for mission critical applications, standardizes data protection, implements disaster recovery, automates routine tasks, and meets SLAs. The goal in this phase is a near 100% virtualized data center.
- **Dynamic.** In this phase, the IT infrastructure is tightly integrated with IT and business processes. As administrators apply security, performance, and availability policies, the virtualization platform responds automatically without manual interaction. This is the really the beginning of a true private cloud.
- **Cloud.** The cloud or final phase provides a real-time consumption model that meets the descriptions and definitions detailed previously. At this phase, business owners only pay for what they consume and can quickly provision and decommission resources as needed. Control shifts into the hands of the application owner, allowing for management of an extremely fluid environment that instantaneously responds to change across distributed resources regardless of whether they are owned or leased from/hosted by a third party. This entire process is completely transparent to the application and its administrators.

This evolutionary process is just beginning. ESG believes that most organizations remain in the early stages of the cycle, caught between the consolidation and enterprise phases. What's holding them back? Many things. Some organizations don't have the IT skills needed to proceed to the enterprise phase and beyond. Many virtualization and cloud technologies are relatively new and immature. Finally, many firms are finding that virtualization, consolidation, web-based applications, and cloud services are a mismatch for their existing network infrastructure. Yes, cloud computing is the eventual destination, but a long journey lies ahead.

The Cloud Depends Upon the Network


Regardless of their plans, timeframe, or maturity, cloud computing is completely synonymous with and dependent upon IP networks as a foundation. Why? Cloud computing includes IT resource consolidation, Web-based applications, and mobile users who access browser-based applications on mobile PCs, PDAs, smart phones, and a potpourri of innovative new devices. As if this weren't challenging enough, cloud computing also highlights Web 2.0 technologies like voice and video that demand high performance/low latency connections. The network will also be essential in cloud migration as large organizations move TBs of internal information to public cloud providers or between private clouds for backup and disaster recovery.

The linkage between the cloud and the network means that large organizations building toward cloud computing must bolster network bandwidth, equipment, and operations. They will also need robust network visibility to assess traffic, application-level performance, and network latency. ESG's research indicates that this network investment is already underway—nearly half of mid-sized (i.e. less than 1,000 employees) and enterprise (i.e. more than 1,000 employees) will increase spending on network hardware in 2010 (see Figure 2).²

² Source: ESG Research Report, [2010 IT Spending Intentions Survey](#), January 2010.

Figure 2. Network Hardware Spending in 2010

To what extent will your organization's spending on network hardware change in 2010? (Percentage of respondents, N = 264)



Source: Enterprise Strategy Group, 2010.

WAN Optimization is a Critical Underpinning of Cloud Computing

Through this transition, business applications will continue to move from branch offices and even corporate data centers to dynamic cloud computing resources that live and move around the Internet. This migration promises operational and cost benefits, but what about application performance to the growing army of remote users? This issue is concerning since today's Web 2.0 applications depend upon "chatty" protocols, multiple application elements per Web page, and low latency connections.

Fortunately, large organizations have already faced a similar set of problems with their current efforts with data center consolidation and web application development. Many CIOs responded with strategic investments in WAN acceleration technologies to connect data centers together or link data centers to branch offices and remote workers. ESG research indicates that this trend will continue: 27% of large and small organization will make "significant investments" in WAN optimization technologies over the next 12 to 18 months.

Today's WAN optimization implementation is based primarily on physical devices. In a typical implementation, organizations deploy a large high performance WAN optimization system in their data center and smaller versions at the edge of each branch office network. This deployment architecture will continue to be an effective way to accelerate application traffic for corporate-based private clouds.


Will Cloud Computing Change WAN Optimization?

WAN optimization offers clear benefits today, so it is logical that this will continue in parallel with the move to cloud computing. That said, will cloud computing drive new WAN optimization services and innovation? Yes. As cloud computing matures over time, WAN optimization technologies must respond with:

- **Virtualization options.** As large organizations move through the enterprise and dynamic phases, they will virtualize applications and not just servers. This means that Web, application, and database servers will be grouped together as virtual applications that run as VMs and move around the cloud infrastructure in unison. This could improve resource utilization, but what happens to end-user performance when a virtual application hosted in New York City moves to Berlin? To address this type of dynamic IT in the future, virtual applications must be supported and accompanied by virtual WAN optimization appliances as they move around the network. Ideally, virtual WAN optimization appliances will "phone home" when they are deployed in the cloud, connect to command-and-control nodes, and inherit existing policies and configurations immediately. These capabilities will provide traditional WAN optimization services and also let IT professionals easily fine-tune application performance.

- **Edge services.** With more and more network-based activities, it may be advantageous to move some existing data center-based network services like DNS, DHCP, and Active Directory domain controllers to the branch office edge. Running these network services locally can help remote users gain network access and find resources regardless of whether they reside on the corporate network or dynamically move around the network cloud.
- **Mobile device coverage.** Cloud computing is evolving in lockstep with increased mobility. More and more workers are telecommuting each day and many mobile workers access business applications with laptop computers, smart phones, and PDAs. Meanwhile, wireless network technology bandwidth continues to increase with the deployment of WiMax and 4G. WAN optimization solutions must accommodate this growing trend as mobile workers need high performance, low latency network connections to business applications whether they are at the corporate headquarters or connecting wirelessly from the road.
- **Additional protocol support.** Leading WAN optimization technology can already accelerate and compress common application and network protocols like CIFS, TCP/IP, HTTP, and SSL. Okay, but cloud computing growth promises an army of new applications, services, and protocols over time. As this occurs, leading WAN optimization vendors with strong R&D resources have a distinct edge as they can develop, test, and roll out new application and protocol support faster than either laggards or startups.
- **Network visibility.** As cloud computing takes hold, applications and low latency services will quickly multiply. This will present a profound challenge to network engineers. Since throwing bandwidth at the problem is counterproductive, cloud computing will demand a much more granular look at network traffic from Layer 2 through Layer 7. How much bandwidth does each application consume? What is the typical behavior of each network service? How are low latency applications affected by traffic spikes? To answer these questions, CIOs will need network traffic management solutions that provide extremely granular visibility and reporting into network utilization—in real time.

Figure 3. Cloud Computing Requirements for WAN Optimization



Source: Enterprise Strategy Group, 2010.

These WAN optimization features will develop over time for the most part as WAN optimization vendors adjust to cloud computing requirements. To get a glimpse of the future, however, look no further than [Riverbed Technology](#), a leading provider of WAN optimization solutions. Most networking professionals know Riverbed for its family of Steelhead appliances for WAN optimization, but Riverbed's product portfolio is actually much broader. Steelhead appliances already provide edge services like those mentioned about through its Riverbed Services Platform (RSP). Furthermore, Riverbed offers Steelhead as a virtual appliance for WAN optimization as large organizations move applications to the cloud. For mobility, Riverbed's Steelhead Mobile Software accelerates network performance directly to user PCs. Finally, Riverbed's Cascade product line provides the company with a leading solution for network and application visibility.

Given its product portfolio, customer base, and market leadership, it is safe to assume that Riverbed will extend its Steelhead solutions and underlying Riverbed Optimization System (RiOS) to support new cloud-specific requirements in the future. As such, CIOs should feel comfortable partnering with Riverbed for both tactical and strategic needs.

The Bigger Truth

ESG believes it is critical that users move beyond cloud computing hype to understand the technology for what it is. The on-demand cloud computing model has the potential to alter the way IT applications and infrastructure are deployed, but it won't happen overnight.

CIOs should be realistic about cloud computing and prudent about the ongoing transition. One way to do this is to assess where your organization sits in ESG's cloud evolution model and then plan for the incremental phases accordingly. Remember that cloud computing can only be built on top of a stable and saleable network foundation. Smart IT executives will assess their network capabilities and business plans to scope out a network infrastructure that delivers performance, flexibility, and scale for future needs. As this paper states, WAN optimization is a key component in this plan, so make sure to choose a WAN optimization solution that supports current and impending cloud computing needs.



Enterprise Strategy Group | **Getting to the bigger truth.**